The First 4 Hours

An Incident Response Field Guide — Volume 1: The Four Hours

Sylvan Press

Sylvan Press

Field-tested, plain-English references for the people who do the work and stand behind their call.

Contents

The First 4 Hours

An Incident Response Field Guide — Volume 1: The Four Hours

Published by Sylvan Press, the security imprint of Sylvan Assurance, LLC — Vermont, USA.

ISBN (paperback / hardcover / ebook): to be assigned prior to publication.

Print and ebook ISBNs are registered through Bowker before release.

First edition, 2026.

Disclaimer. This book provides general guidance and recommended security and compliance practices, not legal, financial, or professional advice, and it does not create any advisory or consulting relationship between the author, the publisher, and the reader. Every recommendation is optional — a widely accepted way to reduce a common risk — and the reader decides which to adopt, adapt, or decline. Following this guidance reduces exposure to common risks but does not guarantee any particular outcome and does not prevent any particular breach, incident, or loss. Where this book refers to laws, regulations, standards, or commercial products, it does so for general awareness only; confirm specific obligations with qualified counsel, and note that product references are illustrative and do not constitute endorsement.

Sylvan Press — sylvanassurance.com

Preface

Incident response has a literature problem. There are shelves of good books on how to *prevent* incidents — security architecture, identity and access management, network segmentation, vendor risk — and very little on what to do in the four hours after one arrives. That gap is where most of the avoidable damage gets done.

This book is about those hours: the ones that begin the moment somebody says, into a phone or a Slack channel or the doorway of an office, *I think we have a problem*. It assumes the prevention work has been done, half-done, or not done at all, because once the alert has fired, none of that changes what to do next. What changes the outcome is whether the response closes off fewer of its future options, preserves more of the evidence the longer investigation will need, and sets off fewer of the second-order problems that turn a manageable incident into an unmanageable one.

The frameworks in this book come from more than two decades of practitioner experience inside a top-25 global company — across network security, corporate information security, computer security incident response, and product security incident response. Across all of it, one pattern holds: the outcome of an incident is set, more often than not, in its opening hours, by people who are tired, under-informed, and making hard-to-reverse decisions on partial facts. The teams that come through cleanly are rarely the ones with the most tooling. They are the ones that have thought about those first hours in advance —

that know what to preserve, what to say, whom to call, and what not to do.

Four recurring organisations carry the worked examples, chosen to span the range of who actually responds to incidents in the small and mid-market world: Halberd Consulting, a one-person practice; Fernwood Supply, a family-owned wholesaler whose security is outsourced to a managed provider; Meridian Health Group, a regional health-services group with one or two security people doing everything at once; and Northridge Federal Bank, a mid-cap bank with a mature function and a 24/7 operations centre. One of them is probably close enough to your situation that you will recognise yourself.

It is worth being honest about what the book does not do. It will not make any organisation incident-proof; that is not a state real programmes reach. The most carefully built programmes still have incidents — they have fewer, catch them earlier, and respond more cleanly, but they have them, and any author or vendor who promises otherwise should be treated with care. What the book offers is a structured way to work the four-hour frame, drawn from real incidents at every scale, so that the next one is smaller, better-evidenced, and more defensible than it would otherwise have been.

These frameworks have been used in practice, not modelled in a seminar. Each decision is described so it can be made at 4:51 on a Friday afternoon, before a long weekend, with half the facts anyone would want. That is the moment this book was written for.

How to Use This Book

A reference book is only as useful as the path the reader takes through it, and the right path depends on why you opened the book in the first place. The chapters are sequenced for a reader who is new to incident response work and is reading the volume straight through to build a mental model, but very few practitioners will read it that way. What follows is a short guide to the paths that hold up in practice, organised by the situation that brought you here.

0.0.1 The shape of this volume

The volume is in three parts, plus a working appendix. The parts are not the same length, and they are not all meant to be read with the same level of attention.

Part I — The First Hour (Chapters 1 through 7). The opening of an incident, minute by minute. Chapter 1 sets the frame. Chapter 2 is about how incidents actually announce themselves — the seven arrival channels, their latencies, and the verification discipline for each — because the phone call is the last mile of detection, not the first. Chapter 3 is the first sixty minutes in detail — what to do, in what order, with what tools, while you are still trying to figure out what is actually happening. Chapters 4 through 7 cover the four pieces of work that you almost always have to do in that first hour: triage without panic, the containment-versus-preservation tradeoff, the holding statement that keeps the internal narrative coherent, and the call tree (who to wake,

and when, and on which number). Read Part I in order, especially if you have never run an incident before. It is the part of the book that builds the muscle memory for the opening.

Part II — Hours One to Four (Chapters 8 through 13). The second through fourth hours, when the picture starts to firm up and the decisions start to bite. Giving the response a room and a rhythm, scoping the blast radius, evidence preservation in real time, customer-facing communications, when the regulatory clock starts, and escalation discipline as the cast of the response grows. Part II is also sequential, but the chapters are largely self-contained — if you already have your evidence-preservation practice locked down, you can skim Chapter 10 and read Chapter 11 (customer communications) more carefully.

Part III — The Operator Under Load (Chapters 14 and 15). Two chapters about the person doing all of the above. Chapter 14 is about hour three — fatigue, stress, and the quality of the decisions you are still capable of making when the adrenaline is gone and the irreversible calls are arriving. Chapter 15 walks the four teams to the four-hour mark and takes honest stock of where you stand: what is stable, what carries forward, and what kind of work comes next.

Appendix A — The Sixty-Minute Triage Checklist. The working companion to Part I: the sequential checklist for the opening hour, in the infrastructure-incident and product-vulnerability variants, with the tier-specific overlays. The orientation in Appendix A walks you through customising it. Customise it, print it, and put your copy where your hands will find it — most operators tape it inside the front cover of the Workbook volume.

0.0.2 The three volumes

This volume covers the four hours themselves. It does not cover everything, and you should know where the rest lives before you need it.

This book — The Four Hours. The universal playbook: detection through escalation, plus the operator under load. Organization-agnostic; everything here applies whether you are one person or a 24/7 security operations centre.

The Four Organisations, and After. The playbook run at four ma-

turity levels — solo operator, MSP-served small business, mid-market with internal IR, enterprise IR function — followed by the specialised incident shapes (ransomware, supply-chain compromise, the vendor's breach that becomes yours, insider and lost-device cases) and the aftermath: the post-incident review, regulatory follow-up after the window closes, and what to change before the next one. The post-incident review template set and the IR Maturity Self-Assessment ride with it as appendices.

Volume 3 — The Operator's Workbook. The artifacts: holding-statement and customer-communication libraries, escalation trees and decision aids, the evidence-preservation runbook, the regulator notification matrix, the template finder, and the IR glossary. Built for the live moment, the way a glove compartment is organised for the breakdown rather than the showroom.

The complete set is under $30 in ebook. If you buy only one volume for the shelf nearest the phone, buy the Workbook; if you buy only one to read first, it is this one.

0.0.3 Reading paths

The chapters sequenced above are the long-form reading path. The paths below are the shorter ones, organised by why you opened the book.

You are standing up an IR practice for the first time, in a small organisation. Begin with Chapter 1 for the framing, and then read Chapter 7 (who to wake, and when) and Chapter 13 (escalation discipline), because those are the two areas in which newly-established IR practices most often have the largest operational gaps. Once those are absorbed, return to Chapters 2 through 6 in order, and you will have accumulated enough structure to handle an incoming call without inventing the process around it on the fly. When you are ready for the starting-from-zero portrait of an organisation your size, *The Four Organisations, and After* opens with it — the solo operator's reality is the most concrete from-nothing chapter in the set, even if you do not consider yourself a solo operator.

You have an IR plan that has not been exercised in a while. Take the IR Maturity Self-Assessment (an appendix of *The Four Organisa-*

tions, and After), and score yourself honestly. The rubric will surface the two or three areas in which the gap between your current state and your target state is the widest, and those are the chapters to read first. Most inherited IR practices exhibit a recognisable failure shape — the call tree is out of date, the holding-statement library does not exist, evidence preservation is an afterthought, or the regulatory-clock awareness is concentrated in one person who is also the person most likely to be on the response. Each of those failure shapes is the subject of a dedicated chapter in this volume.

You are inside a managed-service-provider relationship and you are not sure who does what when an incident actually lands. Read Chapter 7 (who to wake) and Chapter 13 (escalation discipline) for the framework that makes the coordination reliable across incident shapes. The full portrait of the MSP coordination dynamic — the office manager and the provider's responder, working the same incident from two buildings — is the second chapter of *The Four Organisations, and After*, and it is worth reading the moment you finish the two framework chapters here.

You are part of a mid-market security team with two or three people. Read Chapters 2 through 6 in order, and follow with a serious pass through Chapters 8 through 13. The mid-market context is the context the book is most directly written for; the majority of the material applies directly to your situation. *The Four Organisations, and After* has the portrait of a team shaped like yours.

You are inside a mature IR function and you are reading this to calibrate. Skim Parts I and II to see what your function does differently from a less-resourced team, because that contrast is frequently the place where the most useful calibration lesson lives. Pay particular attention to Chapter 13 (escalation discipline), which is the chapter mature-IR practitioners who reviewed drafts most often say they wish they had read earlier in their careers. Then read Chapter 14 — hour-three decision quality is the part of the discipline that no amount of staffing buys you out of.

You are coordinating a single live incident right now. Close this book and open the Workbook volume. Find the decision tree, the call-tree skeleton, or the holding-statement template that matches what

you are doing, and use it — and work your customised sixty-minute checklist if the incident is younger than an hour. Return to this volume afterwards, when the incident is closed and you want to understand why a particular step mattered. The Workbook is constructed for the live moment; this volume is constructed for the reflection that follows.

You are an executive sponsor or a board member who wants to understand what your IR team is actually doing. Read Chapter 1 for the frame, and then read the first four chapters of *The Four Organisations, and After* for the four maturity portraits. You will come away with a working vocabulary for the conversations you have with your security leader, and you will recognise the difference between an IR practice that is appropriately mature for its context and one that is under-resourced for what it is being asked to handle. The remainder of this volume can be skimmed.

0.0.4 The four teams as a maturity lens

A short word on the recurring teams, because they are not only a narrative device.

The four teams — Halberd, Fernwood, Meridian, Northridge — represent four points on the maturity spectrum that the customer-side IR market actually occupies. Solo operator. Managed-service-provider-served small business. Mid-market with internal IR. Enterprise IR function. When the book introduces a technique, you will see how each team uses or adapts it. When a maturity question comes up — *how would a real organisation actually do this?* — the four teams are the answer, given at four different scales.

You will not find a fifth team. The book has been deliberately constrained to four because four covers the realistic range without diluting the texture of any one of them. Locate yourself on the spectrum somewhere between two of the teams — most readers find themselves "between Fernwood and Meridian" or "between Meridian and Northridge" — and use that as your reading lens. The Workbook's templates are tagged by the same four-team rubric, so the reading lens you adopt here carries over directly to which template you pull when you need one.

0.0.5 When to consult the Workbook

The Workbook volume and this one are designed to be used together, not redundantly. Most chapters in this volume name the Workbook artifact they pair with — the call-tree skeleton, the holding-statement library, the regulatory-clock cheat-sheet, the customer-notification template. When a chapter says *see the Workbook*, the Workbook has the template, fillable, with the variant you are most likely to need.

If you read only the narrative volumes and never the Workbook, you will have the theory but not the artifacts. If you use only the Workbook and never the narrative volumes, you will have the artifacts but you will not know when they apply or when they need to be adapted. The set is the product.

0.0.6 A final note

The set's reference appendices are not optional reading. The glossary (in the *Operator's Workbook*) keeps terminology consistent across chapters that touch overlapping topics — *containment*, *isolation*, *preservation*, *quarantine* all have specific meanings here, and the glossary holds the line. The maturity self-assessment (in *The Four Organisations, and After*) is referenced repeatedly in the body of this book, and it is the single most-used artifact readers returned to in pre-publication review.

If you are not in any of the situations above, read Chapter 1 next. It is short, and it sets the frame for everything that follows.

About This Book

The First 4 Hours is a three-book set on incident response operations. This book is the four-hour playbook itself. *The Four Organisations, and After* runs the playbook at four organisational maturity levels, through the specialised incident shapes, and into the aftermath. *The First 4 Hours Operator's Workbook* contains the templates, decision trees, call-tree skeletons, holding-statement libraries, and regulatory-clock references that the narrative volumes describe. Each stands alone; references among them appear throughout, and the set is designed to be used together.

The book is positioned as a practitioner field guide grounded in operational experience. It is not a survey of the academic literature, and it is not a vendor pitch. The frameworks here are drawn from incident response work at a range of organisational scales — from the one-person consultancy to the enterprise security operations centre — and from more than two decades of professional security practice inside a top-25 global company, including product-security incident response.

The book is not legal advice, regulatory advice, or insurance advice. Where the book references regulations — the European Union General Data Protection Regulation (GDPR), the United States Health Insurance Portability and Accountability Act (HIPAA), the United States Securities and Exchange Commission cybersecurity disclosure rules, the European Union Network and Information Systems Directive 2, and others — the references are illustrative of how regulatory clocks

tend to shape the first four hours of an incident. They are not exhaustive, and they are not a substitute for advice from qualified counsel in your jurisdiction. Where the book describes how cyber-insurance carriers tend to handle incident-response notification, the description reflects patterns the author has observed across multiple carriers and engagements; it is not a representation of any specific carrier's policy or practice, and it is not insurance advice. Readers whose situations involve regulatory exposure, contractual obligations, or insurance coverage should obtain specific advice from the relevant qualified professionals.

The book frames its recommendations as practices that reduce risk, not as actions that guarantee outcomes. The phrases *is intended to*, *reduces the likelihood of*, *typically*, and *is a recommended way to* carry that meaning consistently throughout. Following the practices in this book substantially reduces the avoidable damage in the first four hours of an incident; it does not eliminate incidents, and it does not guarantee any particular regulatory, insurance, contractual, or operational outcome. The Sylvan Press house style on defensibility wording is deliberate, and the author has not softened it.

The regulatory and insurance landscape that incident response operates inside continues to evolve. The frameworks in this book reflect the state of practice as of the manuscript's preparation in 2026. Specific obligations, notification windows, and carrier expectations change; readers are encouraged to verify current requirements against authoritative sources at the time of any actual incident. Subsequent editions of this book will revise the frameworks as the landscape changes.

0.0.7 Conventions used in this book

A few conventions are worth flagging up front.

Acronyms are spelled out at first use in each major section rather than only at first use in the book, because readers skim and jump, and a single first-use spell-out on page one is not sufficient for a reference text. The acronyms in the *intentionally excluded* list (AI, IT, URL, and a handful of others that have entered general English) are never expanded. A consolidated glossary appears in the Workbook volume (Volume 3, Appendix B).

The four recurring teams — Halberd Consulting, Fernwood Supply, Meridian Health Group, and Northridge Federal Bank — appear by name throughout. They are fictional. The people inside them — Iris Halberd, Renata Vargas, Diana Reyes, Mia Choi, and the others named in their respective chapters — are fictional. The patterns the teams illustrate are not. Each team's profile and recurring cast are documented in the book's character bible, which is not reproduced in the manuscript but is available on request for editorial reference.

Specific times of day, days of the week, and months appear in scenes throughout the book — 7:51 a.m. on a Tuesday in October, 4:48 p.m. on the Friday before Memorial Day, 2:47 a.m. on a Sunday during a long weekend. The specificity is deliberate. The felt experience of the first four hours depends on the time the incident lands, and an incident at 2:47 a.m. on a Sunday is different in kind from an incident at 10:15 a.m. on a Wednesday, in ways that the methodology has to accommodate.

Single quotation marks appear inside double quotation marks where nested quotes are necessary. The spelling is United States English with a small number of conventional exceptions inherited from international security literature.

Where the book quotes a regulation, a framework, or another published source, the citation appears as a chapter-end note rather than as a page-foot note. The choice keeps the running text uncluttered while preserving the reader's ability to verify or follow up on any specific reference.

Chapter 1 begins on the next page.

1 The Phone Rings

It is 7:51 a.m. on a Tuesday in October, and the phone is on the kitchen counter, and the kitchen counter has a coffee cup on it that is one sip in. The screen says the call is from a client. The client is a forty-person accounting firm two towns over whose managing partner does not call before nine unless there is a reason. The reason, when I answer, is that one of the partners cannot get into her email, and the password-reset link is not arriving, and "the bookkeeper says some weird invoices went out yesterday that nobody wrote."

I have walked in on a lot of these. The detail that tells me what this one is, before the managing partner has finished his second sentence, is the weird invoices. The locked-out account is a symptom. The invoices are the incident. Someone is inside the email, the inbox rules have probably been altered so the partner cannot see replies, and the password-reset email is being captured and deleted before she ever sees it. The clock — one of several clocks, as we will get to — started somewhere between yesterday afternoon and this morning, and I am hearing about it now.

I put the coffee down. I do not say anything dramatic. I say: "Okay. Stay on the phone with me. Do not touch anything yet."

This is the moment the book is about.

Not the breach itself. Not the long tail of forensics and notifications and the eventual lessons-learned writeup. Those have their own books — good ones — and I will point to them. This book is about the four

hours that begin the moment somebody says, into a phone or into a Slack channel or in the doorway of your office, *I think we have a problem*.

The first four hours are different from the rest of the incident in ways that are not always obvious until you have worked a few of them. They are the hours when the decisions you make are easiest to make wrong, and when wrong decisions are hardest to walk back. They are also the hours when the people involved are the least prepared to make those decisions, because they are tired or surprised or both. And they are the hours when the systems those people would normally rely on — email, chat, the wiki, the runbook page they wrote eighteen months ago — may be the systems that have been compromised. The first four hours are a particular kind of work, done in particular conditions, and the work has its own shape. That shape is what this book is for.

1.0.1 Who this book is for

This is a book for the practitioner who is responding to incidents, not the executive who is planning for them.

There are many good books about how to *prevent* incidents. They cover security architecture, identity management, network segmentation, vendor risk, the whole apparatus of controls. They are useful, and a number of them sit on my shelf. They are not this book. If you are building a security programme from scratch, or hardening one, start there.

There are also good books about *governance* — about how to organise a security function, how to report up to a board, how to think about risk as a discipline. They are useful too. They are also not this book.

This book starts at the moment the phone rings. It assumes the prevention work has been done, or has been done partially, or has not been done at all, and that none of that matters right now because the alert has fired and the question on the table is *what do I do in the next four hours that will not make this worse*.

The reader I have in mind takes one of four shapes:

- You are the solo practitioner. You are the founder, the fractional Chief Information Security Officer (CISO), the consultant on retainer, the IT-and-security person at a small firm that is too small to have a security team. The phone rings, and you are it. The walk-

through you will recognise most is the solo chapter that opens Volume 2, set at Halberd Consulting.

- You are the contact inside a small business whose security is outsourced to a managed service provider (MSP). The phone rings, and your first call is to the MSP. The dance between the two of you — what you can do, what they should do, who decides what — is the work. The chapter set at Fernwood Supply is about that dance.
- You are one of two or three security people at a growing mid-market company. You have an incident-response (IR) plan, you have run a tabletop or two, you have a cyber-insurance broker on speed dial, and you are still going to feel undersized when a real incident lands. The chapters set at Meridian Health Group are about that gap between *plan* and *response*.
- You are part of a mature IR function inside a larger organisation. You have a 24/7 security operations centre (SOC), retained forensics, retained outside counsel, and a runbook for almost every shape this thing could take. You will recognise the chapters set at Northridge Federal Bank, and you will also recognise that even at scale, the first four hours do not get easier — they get *different*.

I will return to those four teams throughout the book. They are composites — fictional, but assembled from real organisations I have worked alongside or watched from a respectful distance — and they cover the realistic spread of who responds to incidents in the small and mid-market world. Halberd, Fernwood, Meridian, Northridge. Solo, MSP-served, mid-market, enterprise. When a topic chapter introduces a technique, you will see how each team uses or adapts it. When a maturity question comes up — "*how would a real organisation actually do this?*" — the four teams are the answer at four different scales.

A note on the names: I will not narratively visit any of them in this chapter. The point here is to tell you who they are so you know which lens to read through. The visits begin in Chapter 3.

1.0.2 Why the first four hours are different

Four hours is not a magic number. It is an honest one.

If you have ever read the regulatory literature on incident response — the EU General Data Protection Regulation (GDPR) Article 33

seventy-two-hour clock, the EU Network and Information Systems Directive 2 (NIS2) Article 23 twenty-four-hour early warning, the US Securities and Exchange Commission (SEC) Form 8-K Item 1.05 four-business-day rule, the US Health Insurance Portability and Accountability Act (HIPAA) sixty-day breach-notification rule — you will notice that none of them has a "four-hour" deadline. The four-hour frame is not a statutory window. It is an operational window. It is the window in which the decisions you make set up everything that happens after.

A handful of things tend to be true about the first four hours that are not true about hours five, twelve, or seventy-two.

Evidence is most volatile early. The contents of random-access memory on a compromised machine, the temporary files an attacker dropped and then deleted, the in-flight network connections, the ephemeral session tokens, the partial transcripts of attacker shell-history that survive only until the responsible process exits — these artifacts decay quickly, and the rate of decay accelerates whenever an administrator who does not yet appreciate the forensic implications of their actions interacts with the affected system. A well-meaning systems administrator who reboots an apparently-misbehaving server in hour two, with the entirely reasonable intention of restoring service quickly, destroys evidence that a forensic firm engaged in hour ten will be unable to reconstruct from any other source, which means that the longer-running investigation has to proceed without the answers that the volatile evidence would have provided. The order in which you preserve, contain, and remediate therefore matters most when the evidence is most fragile, and the evidence is most fragile in the first hours, which is to say precisely when responders are least likely to have a documented sequence and most likely to act on instinct.

Decisions are easiest to reverse early. In hour one, you can decide whether to engage a forensics firm, whether to call the cyber-insurance carrier, whether to invoke attorney-client privilege over the response by routing the work through outside counsel, and whether to communicate internally on the corporate Slack workspace or to migrate the response conversation onto a separate channel that has been provisioned specifically for the incident. Each of those decisions, once made, be-

comes correspondingly harder to unmake, because the consequences of the original decision begin to accumulate immediately and to entangle with the next set of decisions on the table. Once you have written incident notes in a non-privileged Slack channel, the notes exist in a non-privileged Slack channel, and they are discoverable in subsequent litigation regardless of how strenuously the organisation later wishes that the notes had been written somewhere subject to privilege. Once you have rebooted the affected workstation to restore service, the memory contents that would have helped a forensics firm understand the attacker's intent are gone, and they are not recoverable through any later effort. The first four hours have an unusually high ratio of consequential decisions to elapsed time, and a correspondingly low ratio of available context to required decisions, which is a hard combination under any circumstances and a particularly hard combination at the beginning of a long day in which the responder is already operating on incomplete sleep and partial information.

Reporting clocks tend to begin at *awareness*, not at *compromise*. This characteristic of the regulatory landscape surprises practitioners who are reading the rules carefully for the first time, because the natural intuition is that an obligation to notify a regulator about a breach would begin at the moment the breach occurred. The GDPR Article 33 clock instead begins when the data controller becomes aware of a personal-data breach, the HIPAA breach-notification clock begins when the covered entity discovers the breach, and the SEC Form 8-K Item 1.05 clock begins when a public registrant determines that an incident has crossed the threshold of materiality. *Awareness*, *discovery*, *determination* — these are events that occur inside your organisation, on a timetable governed by your detection capabilities and your internal escalation discipline, rather than events that occur at the moment the original compromise took place. The regulatory clocks are accordingly timed from the moment you found out, which is a moment that you, the responder, are the person best positioned to record contemporaneously. The first four hours are when the great majority of "found out" moments happen, and the timestamp you write down in those hours is a timestamp you may very well be asked to defend, in writing, two years later, in front of a regulator who is looking at the documenta-

tion you produced and asking whether the response was timely. Write the timestamp down on paper. Write it down somewhere you will not lose it, somewhere that will not be overwritten by an enthusiastic colleague editing the same document, and somewhere that you can produce later if you are asked to demonstrate when awareness arose.

Communications go off the rails fastest in the first four hours. Word spreads through an organisation under stress the way word spreads through any organisation under stress, which is to say unevenly, with embellishment, through informal channels that the formal incident-response process never anticipated needing to govern. By hour five, somebody who heard from somebody who happened to be in the hallway when the original responders convened has already mentioned something — vaguely, with caveats, with the best of intentions — to a customer-facing colleague who in turn mentions it to a customer. By hour twelve, a regional sales lead has emailed a partner organisation a sentence that includes the phrase "we are seeing some unusual activity," and once the sentence has been emailed, it cannot be retrieved. The first four hours are when the communications discipline you set — who is authorised to speak, on which channel, in which language, with which approvals, and to which audiences — either holds together as a coherent regime or begins to fail in the small ways that compound across the rest of the incident. Most of the avoidable damage I have personally observed in incidents was not done by the attackers responsible for the original compromise; it was done, instead, by well-intentioned people inside the responding organisation who said true things to the wrong audience at the wrong time, often because nobody had taken the trouble to tell them not to.

The team is incomplete in the first four hours. Even at organisations whose incident-response function is mature and well-resourced, the full responding cast is rarely all in the room by minute thirty, and the cast that is in the room is typically a partial subset of the cast that will eventually be required. The technical responders are usually first to arrive, because they were the ones already in front of a console when the alert fired, or because they were the on-call engineer who was paged at 2:47 a.m. and who took the call seriously enough to drive in. Counsel arrives next, often by phone, often from a personal mobile

on which counsel happens to be reachable at the time. The executive sponsor and the communications lead come into the response over the course of the first couple of hours, typically through a sequence of escalations that the responding team has to manage in parallel with the response itself. The cyber-insurance carrier's incident-response hotline, the retained forensics firm if there is one, the outside counsel if the engagement is being routed through them — each of those parties is a phone call that takes time to initiate and additional time to bring up to a useful level of context. The first four hours is therefore the window in which the responding team is assembling itself while also doing the response work, and the two activities interfere with each other in ways that the chapters that follow are partly about: how to assemble the team without slowing down the response, and how to keep the response from running ahead of the team in ways that the rest of the team will later have to undo.

You are tired or surprised, or both. A surprising proportion of incidents land at 7:51 a.m. on a Tuesday, or at 4:48 p.m. on a Friday before a long weekend, or in the middle of the night on a holiday during which the on-call rotation was understaffed because nobody wanted to draw the short straw on Thanksgiving. Attackers do not respect business hours, and the moment of awareness, on your side, is generally whenever you happened to notice something out of the ordinary, which means the first four hours of response are very often the first four hours of a long, badly-timed day in which nothing else on your calendar will be done. The methodology in this book is built around the assumption that you, the responder, are not at your sharpest when the response begins, because the response begins when it begins, regardless of whether you are ready. It is built to be usable when you are running on coffee and one sip of it, or on no sleep, or on adrenaline that is starting to fray at the edges. It is built to be usable, in fact, when you are the only person on the response, and the only checklist you have within arm's reach is the laminated card in the glovebox of your car that you put there six months ago because you thought you might need it some day.

That is what makes the first four hours their own discipline. They are not the rest of the incident, run faster. They are a separate piece of

work, with their own constraints, their own decisions, and their own failure modes.

1.0.3 The shape of a typical opening

Every incident is its own creature, but the openings tend to rhyme. The shape of a typical opening — what happens in the room, on the phone, on the chat thread, in the first thirty to ninety minutes — looks something like the following sequence, regardless of whether the responding organisation is a one-person consultancy or a regional bank with a fourteen-person incident-response function. The vocabulary shifts; the pacing shifts; the participating cast shifts. The underlying sequence is largely the same.

There is, first, the *triggering signal* — the alert, the user report, the customer escalation, the regulator's letter, the journalist's email. The signal arrives through a channel that was not necessarily designed to carry it. A phishing-incident report arrives on a personal mobile number because the user could not get into corporate email. A ransomware alert fires in an endpoint-detection-and-response (EDR) console at 2:47 a.m., and a tier-one security operations centre (SOC) analyst — who has, until that moment, been working a quiet shift — escalates to an on-call engineer who is asleep. A vendor with whom the organisation has a managed-file-transfer integration sends a terse advisory by email at 4:48 p.m. on a Friday, and the recipient, who is already mentally on the train home, reads it twice and decides to call somebody. None of these openings is ideal. None of them is unusual.

There is, second, the *first interpretation*. Whoever received the signal makes an initial guess about what it means. The first interpretation is almost always wrong in some particular — sometimes wrong in scope (the responder thinks one account is compromised; in fact six are), sometimes wrong in nature (the responder thinks it is a phishing campaign; in fact it is a credential-stuffing attack against a different surface), sometimes wrong in timing (the responder thinks the signal is from this morning; in fact the underlying activity has been ongoing for eleven days). The first interpretation is nonetheless a useful artifact, because it organises the next set of actions. It also has a tendency to anchor the responding team for longer than is helpful. Good

incident-response discipline holds the first interpretation loosely and updates it visibly as new information arrives.

There is, third, the *small huddle*. Two or three people, sometimes more, sometimes only one person on a call with somebody who is not even part of the responding organisation, gather to decide what to do next. The huddle is often informal. It may take place in a hallway, in a parking lot, on a personal cell-phone call between people who would normally meet by video, on a side-channel chat that somebody set up because they were not yet sure whether the corporate channel was safe to use. The small huddle is when the response either acquires structure — a designated decision-maker, a single channel of record, a written timestamp, a short list of next actions — or fails to acquire structure and continues as a loose collection of well-intentioned people doing parallel work that does not add up.

There is, fourth, the *first outward call*. At some point in the first hour, somebody picks up the phone to a party outside the small huddle. The party is typically the cyber-insurance carrier's incident-response hotline, occasionally outside counsel, occasionally a retained forensics firm, occasionally a peer organisation that the responding team has agreed in advance to consult. The first outward call is consequential out of proportion to its content. The act of picking up the phone moves the incident from internal-suspicion to externally-acknowledged status, which has implications for privilege, for insurance, for documentation discipline, and for the institutional seriousness of the response from that point forward. Responders who delay the first outward call past hour two usually wish, in retrospect, that they had made it sooner.

There is, fifth, the *deliberate slowdown*. Roughly forty-five minutes in, the responding team — if it is being honest with itself — pauses for thirty seconds to ask whether the pace of the response is sustainable. Almost always the answer is no, the pace is not sustainable, and the team adjusts: somebody is designated as the note-taker, somebody else as the calls-and-coordination owner, the technical work is delegated to a single person rather than three, and the next status check is set for a specific time. The deliberate slowdown is one of the more counter-intuitive moves in incident response, and one of the most valuable. The instinct of any honest responder, in the middle of a live incident, is

to keep moving faster. The work, almost always, gets better when the team slows down for thirty seconds at the forty-five-minute mark.

There is, sixth, the *first decision the responding team would rather not make*. This decision varies in shape — it might be whether to take a production system offline, whether to notify a customer before the picture is complete, whether to file an early-warning regulatory notification before the legal team is comfortable with the wording, whether to call the executive sponsor at home on a Sunday — and the timing varies, but the decision arrives. Somewhere in the second or third hour, the team has to make a call that nobody on the team wants to be the person who made. The methodology in this book is partly about how to make those decisions in a way that is defensible later, and partly about how to make them with enough of a paper trail that they do not become the decisions everybody points to in the post-incident review six months later.

Those six beats — triggering signal, first interpretation, small huddle, first outward call, deliberate slowdown, first uncomfortable decision — do not exhaust everything that happens in a typical opening, but they form the underlying rhythm. When the rhythm is recognisable, the response work has texture and direction. When the rhythm collapses — when, for example, the small huddle never forms, or the first outward call gets delayed past hour three, or the deliberate slowdown never happens and the team is still going at sprint pace at hour two — the response work becomes correspondingly harder to recover. Most of the chapters that follow are organised around those beats, with the underlying assumption that recognising them as they are happening is more than half the battle.

1.0.4 What this book promises, and what it does not

I want to be clear about what you should expect.

This book is intended to make your first four hours more effective. It is not intended to make you incident-proof. Nothing makes anyone incident-proof. The most carefully-built security programmes I have ever worked alongside still have incidents — they have fewer of them, they catch them earlier, and they respond more cleanly, but they have them. Treat any author or vendor who promises otherwise with care.

What this book offers, instead, is a structured way to think about the first four hours, drawn from a couple of decades of doing the work and watching others do it. Specifically:

- A felt sense of *what the work is*. The opening chapters are about the texture of an incident-response opening: what it actually looks like in the room, on the phone, on the chat thread. If you have never run an incident, this is the orientation you need. If you have run several, it is the calibration you may not have had a chance to do, because the people you would calibrate with are also doing their own incidents, and nobody has time.
- A framework for how incidents announce themselves and the first sixty minutes (Chapters 2 and 3), the hours that follow — triage, containment, the holding statement, the call list (Chapters 4 through 7) — and the third and fourth hours (Chapters 14 and 15). The framework is not a recipe. It is a sequence of decision-shaped questions, with the techniques that tend to work for each.
- The topical chapters that practitioners reach for again and again: containment without destroying evidence (Chapter 5); scoping the blast radius (Chapter 9); preserving volatile evidence (Chapter 10); the communications discipline (Chapters 6 and 11); regulatory clocks at a glance (Chapter 12); escalation discipline (Chapter 13).
- Four scaled walk-throughs — one for each of the four teams — showing what the first four hours actually look like at solo, MSP-served, mid-market, and enterprise maturity. Those are the work of this guide's second volume, *The Four Organisations, and After* (Chapters 1–4), where the methodology meets the reality of who you actually are when the phone rings.
- The harder cases — ransomware in the first four hours, supply-chain compromise, the breach that turns out to be your vendor's, and insider and lost-device incidents — also ride in Volume 2 (Chapters 5–8), with the post-incident review that follows them.

The promise is modest by design. If you read this book, and you keep it within reach when an incident lands, and you work the four-hour frame, you will be more effective than you would be otherwise. You will make fewer of the decisions that close off your future options.

You will preserve more of the evidence that lets the longer response work. You will trigger fewer of the avoidable second-order problems — the leaked communications, the missed regulatory windows, the unauthorized public statements — that make a manageable incident into an unmanageable one.

You will not, I want to repeat, be incident-proof. Nobody is. The goal is to be incident-*ready*, in the specific sense that when the phone rings, you have a frame to put around the next four hours that is better than the one you would invent under pressure.

1.0.5 A word on the voice

I am writing in first person throughout the chapters that follow because incident response is, fundamentally, first-person work — something a practitioner does in the moment, with whatever combination of tools, colleagues, vendor relationships, and accumulated judgement happens to be available at the time the response begins. I have spent twenty-six years in this profession, eight of them in product-security incident response and the prior eighteen in network security, corporate information security, and computer security incident response teams (CSIRTs) at one of the top twenty-five global technology companies. I have been the responder on a live incident at 3 a.m. on a Sunday morning. I have been the manager of responders during incidents that lasted weeks. I have been the person on the phone at 3 a.m. with somebody else's responder, talking them through a containment decision that they were nervous about making and that, in retrospect, was the right one. I have made enough of the mistakes catalogued in this book — the rebooted workstation, the leaked communication, the delayed call to counsel, the timestamp that nobody wrote down — that I can describe them honestly, without the kind of distance that turns advice into abstraction.

I also do not want to oversell the experience. Twenty-six years in the profession is enough to have observed patterns across many incidents and many organisations, but it is not enough to have seen everything that incident response can throw at a responder, and any practitioner who claims otherwise should be regarded with care. The four teams in this book are composites of organisations I have worked with, along-

side, and against over those twenty-six years, and the patterns they exhibit are patterns I have personally observed in the wild — but every real incident has its own texture and its own surprises, and the methodology in this book is intended as a starting frame for the responder's thinking, rather than as a script to be followed mechanically. Where I am confident in a recommendation, I will sound confident. Where the evidence is genuinely mixed, or where the situation typically depends on contextual factors that I cannot anticipate from the page, I will say so explicitly, and I will try to give you the tools to make a contextual judgement rather than a generic one. Hedging in security writing is sometimes a stylistic tic that obscures rather than clarifies; here, when you see hedging, the hedging is doing intentional work, and the work is to remind you that incident response is judgment-laden rather than rule-bound.

A few specific choices in the writing:

- I will use specific times (7:51 a.m., 4:48 p.m., the third Tuesday in October) in scenes, because the felt experience of the first four hours depends on the time of day and the day of the week.
- I will name the four teams — Halberd, Fernwood, Meridian, Northridge — by name throughout. They are fictional. The people inside them are fictional. The patterns they illustrate are not.
- I will sometimes address you, the reader, directly, in the voice of somebody who has watched a lot of timestamps get lost in chat scrollback.
- I will avoid war metaphors where I can. Incident response is closer to triage in an emergency room — careful, sequenced, mostly calm — and the metaphors that fit that work tend to fit this work too.

1.0.6 The shape of the next four hours

A useful frame for the four hours: each hour has its own centre of gravity.

The first hour is about *not making it worse*. That sounds modest. It is not. Most of the avoidable damage in an incident is done in the first sixty minutes by responders who act decisively before they have the picture. The work of the first hour is to capture the timestamp, get to a clean channel, contain without destroying, name a single decision-

maker, and start the calls — to the insurance carrier, to counsel, to the forensics firm if one is needed. By the end of hour one, you should have written down five things: when you became aware, what the apparent scope is, which clocks are running, who is engaged, and what the next decision is and when it must be made. Chapter 3 is about hour one in detail.

The second hour is about *getting the picture*. By the end of hour two, the apparent scope is firmer, the data classification is at least drafted, the forensics firm (if engaged) is starting to work, and the communications discipline is in place. The second hour is also when the team is fullest — most of the people who are going to be in the response have arrived or are on the phone — and when coordination overhead is highest. Chapter 4 covers hour two.

The third hour is about *making the first hard decision*. Almost every incident produces a decision in the third hour that cannot be deferred any further. Do we notify the affected customer? Do we shut down the production system to contain? Do we file the early-warning regulatory notification? The decision shape varies, but the third hour is when "we need more time" stops being the right answer to at least one question. Chapter 5 is about that decision and the structures that make it more reliable.

The fourth hour is about *setting up the next twenty*. By the end of hour four, the response is no longer an opening; it is a sustained operation. Shifts get organised. The communications cadence is established. The legal and regulatory clocks are tracked in a single place. The handoff to a longer-running incident-response process is built. Chapter 6 covers the transition.

After Chapter 6, the book moves to topical chapters that you can read in any order, and finally to the four scaled walk-throughs.

1.0.7 Back to the phone

I want to come back, briefly, to the moment I started with. The phone on the counter. The coffee cup with one sip out of it. The managing partner of a forty-person accounting firm telling me about a partner who cannot get into her email and a bookkeeper who is worried about some weird invoices.

What I do next, on that call, is the work of Chapter 3 — the first sixty minutes. I will not narrate it here. I will say only this: my first words, after *do not touch anything yet*, are to ask the managing partner to step away from his own laptop, find a notepad and a pen, and write down the current time. Then I ask him to tell me everything he knows, slowly, while I make notes on my own paper. I do not open a Slack thread. I do not start a Zoom. I do not email anyone yet. The two of us, on the phone, with paper, are the entire incident-response team for the next eleven minutes, until the partner with the locked-out account joins on a different phone from a different building, and then there are three of us. By minute thirty, there will be five. By the end of hour one, the cyber-insurance carrier will be on the line, a forensics firm will be inbound, and counsel will have called back. That is the work of the chapters ahead — beginning with how the signal arrives at all, which is Chapter 2's subject, and then the first sixty minutes in Chapter 3.

The point of telling you about the phone on the counter is that the work begins exactly where I described it. Not in a war room. Not at a console. On a kitchen counter, with a phone in one hand, with a managing partner on the other end who does not know what he does not know, with a coffee cup cooling down and the rest of a Tuesday stretched out ahead. That is the felt experience of the first four hours, and that is what we are going to spend the rest of the book getting good at.

Recap: The first four hours of an incident are a distinct kind of work — short, consequential, evidence-fragile, communications-fragile, and almost always under-resourced. This book is for the practitioner doing that work, not the executive planning for it. Four recurring teams — Halberd Consulting (solo), Fernwood Supply (MSP-served), Meridian Health Group (mid-market), and Northridge Federal Bank (enterprise) — will appear throughout to show the methodology at four different scales.

Next: Chapter 2 — Detection Sources and the First Signal. How incidents actually announce themselves — and why the four-hour clock you can control usually starts long after the incident did. Then Chapter 3 — The First Sixty Minutes: the timestamp, the clean channel, the single decision-maker,

the network-isolate-don't-power-off rule, and the five things you want written down before minute sixty-one.

You've just read Chapter 1 of

The First 4 Hours

The full book runs 330 pages across 15 chapters,
with a complete back-of-book index.

Published by Sylvan Press, the book imprint of Sylvan Assurance —
field-tested, plain-English references for the people who do the work.

Scan to see the full book — and where to buy it.
sylvanassurance.com/go/first-4-hours

Not ready for the full book? See where you stand first — free, a few minutes:

sylvanassurance.com/go/free-first-4-hours