

Cloud Security

*A Practitioner's Field Guide — Volume 1:
Foundations and Control Domains*



Sylvan Press

SYLVAN PRESS

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

Contents

Preface	ii
How to Use This Book	iv
Reading Paths	v
The Four Organisations	vii
When to Consult a Professional	viii
A Final Note Before You Begin	ix
About This Book	x
1 What Cloud Security Actually Is	1
2 The Shared-Responsibility Model in Practice	14
3 Cloud Asset Inventory	32
4 A Cloud-Security Maturity Spectrum	50
5 How Cloud Security Differs from On-Prem Security	67
6 Identity and Access Management at Cloud Scale	84

CONTENTS

7	Secrets, Keys, and the Cryptographic Surface	102
8	Network and Data-Plane Security in the Cloud	121
9	Configuration Baselines and Drift Detection	139
10	Logging, Monitoring, and SIEM Integration	158
11	Data Security: Classification, Encryption, and Residency	175
12	Workload Protection: Containers, Serverless, and VMs	193
13	Kubernetes and Container-Orchestration Security	212
14	Infrastructure as Code Security	227
15	SaaS Security Posture	241
16	Cloud Cryptography Patterns	254
17	Edge, CDN, and API Gateway Security	267
18	The Platform Secured, and What Comes Next	280
A	Appendix A	292
	A.1 The Control-Surface Artefacts Quick Reference	292
B	Appendix B	298
	B.1 Acronyms in Plain English	298
	Index	308
	About Sylvan Press	312

Cloud Security

A Practitioner's Field Guide — Volume 1: Foundations and Control Domains

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

Copyright © 2026 Sylvan Assurance, LLC. All rights reserved.

Licensed for the reader's own personal and organisational use. No part of this publication may be reproduced, distributed, or transmitted for resale or redistribution without the prior written permission of the publisher.

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

Cloud security fails in the gap between the architecture diagram and the console. The diagram shows a clean topology with the right controls in the right places; the console shows the account somebody spun up last Tuesday with a public bucket and an over-scoped role. This book is about closing that gap with controls that hold — domain by domain, account by account, at the pace accounts actually get created.

The control domains, architectural patterns, and failure modes in this book come from more than two decades of practitioner experience inside a top-25 global company, observed from alongside cloud security teams and from the incident-response and product-security side of cloud problems — next to the cloud function rather than inside it, under real conditions rather than in a reference diagram. The recurring lesson is that cloud security is less about exotic attacks than about ordinary controls applied consistently across an estate that no single person can see end to end.

This book builds the foundations — what cloud security actually is, the shared-responsibility model as it works in practice rather than on the slide, asset inventory, the maturity spectrum, and the real differences from on-premises practice — and then walks the control domains: identity and access management at cloud scale, secrets and the cryptographic surface, network and data-plane security, configuration baselines and drift detection, logging and monitoring, data security, and workload protection.

Four recurring organisations carry the worked examples: Stagheart, Helix, Westmoor, and Brevenmark. Each runs its cloud estate at a different scale and maturity, and a control that is right for one is wrong for another. One of them is probably close enough to your own situation that you will recognise yourself.

A word on what the book promises. The practices here are framed as approaches that reduce risk and have been seen to work in real environments — not as actions that guarantee any outcome. The aim is an estate you can defend on the ordinary Tuesday, before the account that no one remembers creating becomes the account in the incident report.

How to Use This Book

0.0.1 The shape of this volume

Part I — Foundations (Chapters 1 through 5). What cloud security is, shared responsibility in practice, asset inventory, the maturity spectrum, and how cloud differs from on-prem. Read in order if the domain is new to you.

Part II — Control Domains (Chapters 6 through 12). IAM, secrets and keys, network and data plane, baselines and drift, logging and SIEM, data security, and workload protection. Self-contained; read what your estate exposes first.

Part III — The Platform Layer (Chapters 13 through 18). Kubernetes and container orchestration, infrastructure as code, SaaS posture, cloud cryptography patterns, and the edge.

The Cloud Security set. This book is complete on its own: the foundations, the control domains, and the platform layer — the full security architecture. Two companion books extend the practice for readers who want them. *Operations and Incident Response* covers vendor diligence, cloud IR and forensics, tabletops, multi-cloud and hybrid coordination, and regulated workloads. *The Cloud Security Operator's Workbook* holds the templates, with the Cloud Security Glossary and the Maturity Rubric. Each stands alone; none assumes you have read the others. The complete set is under \$30 in ebook.

Reading Paths

Three paths through the book fit three situations, and naming them may save you a month.

The newcomer with the questionnaire. Something has just made cloud security your problem — a client’s vendor questionnaire, an auditor’s first cloud question, a board member who read a headline. Read Part I in order, because it was built for exactly this arrival: what you own (Chapters 1 and 2), what you have (Chapter 3), where you honestly stand (Chapter 4), and what carries over from any security you already know (Chapter 5). Then read Chapter 6, because identity is where the next failure is most likely living, and return for the remaining control domains as your estate demands them.

The practitioner transferring in. You have run security somewhere — a network, a SOC, an on-premises estate — and the cloud is the part you have not owned before. Start with Chapter 5, which maps your existing craft onto the cloud directly, then Chapter 2 for where the responsibility line actually sits, then go straight at the control domains in Part II in whatever order your first assignment dictates. The platform layer (Part III) will read quickly for you; its job is to name the surfaces — orchestration, infrastructure code, SaaS, cryptographic defaults, the edge — that fall outside every traditional remit, which is precisely why they fail.

The owner of an inherited estate. The cloud is already there — built by predecessors, grown by teams, never assessed as a whole —

and your question is where to look first. Read Chapter 3 and run the inventory before anything else, because every other chapter assumes you know what you have. Then Chapter 4 to place the estate honestly on the maturity spectrum, then the closing chapter's interrogative list, which turns the whole book into the audit you are about to perform. The chapters the audit flags will name themselves.

The Four Organisations

The examples run on four recurring organisations spanning the size spectrum — a fifteen-person product-design agency with one AWS account, a six-hundred-person cloud-native mobile operator paying down six years of speed, a research university whose federated estate no one can mandate, and a fifty-thousand-person multinational insurer running a platform-grade programme across three providers. They are introduced properly in Chapters 1 and 4, and they exist so that every discipline in the book appears at four price points: when a chapter's full apparatus would be over-engineering at your scale, the smaller organisations show the version that is not, and the lesson is always proportion rather than imitation. They are composites, assembled from observed patterns; any resemblance to specific organisations is the genre, not a report.

When to Consult a Professional

This book describes recommended practices for securing cloud environments; it does not know your estate, your contracts, or your obligations. Three situations call for help beyond any book. Where decisions touch legal or regulatory exposure — international data transfers, breach notification, sector rules such as NIS2 or DORA, contractual security commitments — qualified counsel should see the specifics, because the landscape moves and the penalties are real. Where a live incident is unfolding, incident response expertise, in-house or retained, takes precedence: this book builds the architecture that prevents and survives bad days, and its companion covers operations, but neither substitutes for the people who work them. And where the cloud carries workloads whose failure has physical or safety consequences, the relevant engineering and safety authorities outrank every prioritisation framework here.

A Final Note Before You Begin

The work in this book is not glamorous. It is account lists and key rotations, tagging policies and log destinations, quarterly reviews that find nothing — until the one that finds something. That is the point. The cloud estates that fail rarely fail for want of sophistication; they fail because a surface nobody owned drifted, quietly, until an ordinary threat found it. The chapters ahead are an argument that the unglamorous version, owned explicitly and run on a cadence, is what cloud security actually looks like from the inside — and that it is within reach of a fifteen-person agency, a stretched platform team, and everyone larger. Own the surfaces. The rest is calibration.

About This Book

A few conventions are worth flagging. Acronyms are spelled out at first use in each major section, not only at first use in the book. Where the book recommends a practice, the recommendation is framed as an approach that reduces risk — not as a guarantee of any outcome. Nothing in this book is legal advice; regulatory and contractual specifics belong with qualified counsel, and the landscape moves — verify current requirements against authoritative sources.

Chapter 1 begins on the next page.

1 What Cloud Security Actually Is

Esme Caldwell found the problem on a Tuesday in March, two weeks before the pitch meant to close Stagheart Studio's largest retainer yet. The prospective client — an 80-person FinTech, Series B just closed — had attached a vendor security questionnaire as a condition of signing. Esme had answered a few of these before, mostly by describing what Google Workspace did and hoping the question was about email. This one was different.

Question 34 asked whether Stagheart enforced least-privilege access controls in its cloud infrastructure, and whether it could show its IAM policy review cycle. Esme opened the AWS console — which she logged into about once a month, to check the bill. Eight access keys. She did not recognise three of the usernames. The oldest, created in 2022, had last been used six weeks earlier. She did not know by whom, or for what.

She went looking for a way to answer Question 34 with confidence. What she found was the shared-responsibility model: the provider secures the cloud, and the customer secures what runs in it. The customer. That was her. It had always been her. Nobody had told her that was the deal.

The FinTech signed two months later. In between, Esme cut eight

access keys to four, each with a named owner; turned on CloudTrail in every region, not just eu-west-1; and answered Question 34 honestly — current practice, gaps named, remediation plan attached. The client's security team said it was the most straightforward response they had had from a vendor Stagheart's size. Not because it was perfect, but because it was true, and it showed she knew what she owned.

Esme did not fail because she was careless. She failed because nobody had explained what she owned — and the keys were not all of it: a client-briefing bucket had been public for fourteen months, with no logs to say whether anyone had opened it. The gap between what cloud customers believe they own and what they actually own is where most cloud security incidents begin.

That gap is the through-line of this book, and the thesis is worth stating plainly: cloud security is not a new discipline but a specialisation of the one you already practise — asset inventory, access control, logging, incident response, vendor due diligence — extended to infrastructure that someone else operates. This chapter sets out the working definition and the boundaries that matter; the four organisations whose arcs run through the book arrive at the end of it. How Esme built that honest questionnaire answer — why naming your gaps beats a confident-sounding claim — is worked through in the *Cloud Security Operator's Workbook*.

1.0.1 The working definition

Cloud security is the application of security disciplines to workloads, data, identities, and services that are operated, in whole or in part, on infrastructure provided and managed by a third-party cloud provider.

That definition is more specific than it might first appear, and each clause is doing work.

Workloads, data, identities, and services. Cloud security is not one programme. It is at minimum four: a programme for protecting what you run (workloads), a programme for protecting what you store and process (data), a programme for protecting who can do what and to what (identities), and a programme for protecting the services you expose and consume (APIs, SaaS connections, microservice meshes). In prac-

tice these four are deeply entangled. A misconfigured identity leads to a workload being hijacked, which exfiltrates data through a service boundary. The entanglement does not mean you can have one programme that calls itself “cloud security” and calls it done. It means the four programmes need to be coordinated, share a common asset inventory, and escalate incidents through the same chain.

In whole or in part. Very few organisations are entirely on-premises or entirely in the cloud. Most are in a hybrid state — some workloads in a traditional data centre, some in one or more cloud providers, some in SaaS that sits beneath both. Cloud security applies wherever the cloud portion exists, which in 2026 is nearly everywhere. The “in part” clause matters because it prevents the category error of organisations assuming that because their core ERP is on-premises, they have limited cloud exposure. The SaaS identity provider, the cloud storage bucket for backups, the development environment spinning instances on a personal credit card — these are cloud exposure, regardless of where the production system lives.

Operated, in whole or in part, on infrastructure provided and managed by a third-party cloud provider. This is the clause that does the most organisational work. The cloud provider manages the infrastructure. The customer does not. The servers, the network gear, the physical data-centre perimeter, the hypervisor layer — these are the provider’s responsibility in the shared-responsibility model. This does not eliminate the customer’s security responsibilities. It relocates them. The customer now owns the security of what runs on the infrastructure, what is stored on it, who can access it, and how it is configured. The customer does not own the security of the infrastructure itself — and this is the fundamental adjustment that organisations like Stagheart have to make.

1.0.2 What the shared-responsibility model actually says

The shared-responsibility model is the foundational concept of cloud security, and it is the concept most frequently misapplied. Every major cloud provider publishes a version of it. The details differ across providers, across service types — Infrastructure as a Service (IaaS), Platform as a Service (PaaS), Software as a Service (SaaS) — and

across individual service configurations. The core logic holds across all of them.

The provider is responsible for securing the physical infrastructure: data centres, power, cooling, networking hardware, the hypervisor, and the global backbone that stitches regions together. The provider's security at this layer is, in general, substantially better than what most organisations could afford to operate for themselves. This is one of the genuine and underappreciated security advantages of cloud adoption. For organisations that were running on-premises in a shared colocation facility with a single-factor VPN and a firewall that had not been patched since the last consultant left, moving to a major cloud provider almost always improves the security of the infrastructure layer.

What the provider does not secure — what the customer secures, in every mainstream cloud model — is the layer above. The customer configures access controls for the services they run. The customer decides who has IAM permissions and at what scope. The customer decides whether an S3 bucket is private or public. The customer decides whether database instances are reachable from the internet or isolated in a private subnet. The customer decides whether CloudTrail or equivalent audit logging is on or off, and in which regions. The customer decides how secrets are managed, whether encryption is enabled at rest and in transit, and whether the instances they run are patched. The customer decides whether the configurations they have made drift from the secure baseline they intended.

The misunderstanding Esme had — the misunderstanding most organisations have when they first encounter the cloud — is the assumption that because they are renting infrastructure from a security-conscious provider, they have outsourced the security of their systems. They have not. They have outsourced the security of the infrastructure those systems run on. The security of the systems themselves remains entirely theirs.

The practical consequence of this misunderstanding is not minor. When major cloud providers report that the vast majority of cloud security incidents involve customer-owned configurations rather than provider infrastructure failures, they are not defending themselves against criticism. They are accurately describing where the attack sur-

face lies. The bucket that is public because a developer set it public. The database that is internet-facing because a firewall rule was too broad. The root account that has no MFA because nobody got around to enabling it. The access key that belongs to an ex-employee because the offboarding process did not include IAM. These are customer failures in the customer's layer. The cloud provider's security controls at the infrastructure layer did not and could not prevent them.

The first discipline of cloud security is therefore not a technical one. It is an organisational one: ensure that everyone in the organisation who makes decisions about cloud services — every developer, every infrastructure engineer, every product manager who spins up a new environment — understands what the shared-responsibility split means for the decisions they are making. That is harder than it sounds. The cloud console makes it very easy to deploy something and very easy to misconfigure it. The provider's defaults are sometimes, but not always, the secure defaults. And the feedback loop is slow: a misconfigured bucket may sit public for months before a security scanner, a researcher, or an attacker notices.

The second discipline follows from the first: the organisation needs a mechanism to observe its own configuration continuously, not just at the moment of deployment. Static audits and annual reviews are insufficient in an environment where infrastructure can be provisioned, reconfigured, and destroyed in minutes by any engineer with a console or a Terraform repository. Cloud security posture management (CSPM) tools exist to address this — Chapter 9 covers their use in depth — but the tooling is only valuable if someone is assigned to triage the findings and a process exists to close them. Tooling without process produces an alert queue that nobody is responsible for, which produces a false sense of coverage.

1.0.3 The three ways organisations encounter cloud security

Organisations encounter cloud security in three distinct phases, and the phase largely determines what the first priority should be.

Phase 1: Cloud adoption without cloud security discipline. This is Stagheart. The organisation has adopted cloud services — infrastructure, SaaS, PaaS, some combination — without building the corresponding security disciplines. Access keys are long-lived and informally managed. There is no asset inventory beyond institutional memory. Logging is incomplete or absent. Configuration baselines do not exist. Vendors were procured on their product merits; their security posture was not assessed. The organisation is often unaware of the exposure it has accumulated because it has no mechanism to observe it. Incidents in this phase tend to be discoveries rather than detections: the bucket that turns out to have been public, the key that turns out to belong to a leaver, the SaaS tenant with no offboarding process.

The first priority in Phase 1 is not a sophisticated security programme. It is a set of minimum-viable controls that close the most exposed attack surface. A complete asset inventory of cloud services in use. Multi-factor authentication (MFA) on all privileged accounts and ideally on all accounts. Rotation and expiry of long-lived access keys. Audit logging enabled across all regions. At least one person with a clear mandate to own these controls. That person does not need to be a cloud security specialist. They need to be someone who understands what the controls are for and who has the organisational standing to enforce them.

Phase 2: Cloud-native growth outpacing cloud security. This is Helix Mobile. The organisation has been building on the cloud from day one, or has moved significant workloads to the cloud and kept building. The problem is not ignorance — the engineering teams know their way around the platform — but accumulation. The IAM estate has grown organically and has never been rationalised. There are long-lived keys because they were convenient and nobody went back to retire them. The network is flatter than it should be. Logging is on but nobody watches it. A CSPM tool was enabled at the free tier and never tuned. The organisation is aware that cloud security is a gap; the question is which gaps to close first and in what order.

The first priority in Phase 2 is triage: identify the highest-risk exposure in the existing estate and close it before extending the programme further. For most organisations at this phase, the answer is IAM. An

IAM estate with several hundred long-lived access keys, inconsistent role assumption, and an overly permissive managed policy attached to engineering teams during the startup phase and never removed is a single compromised credential away from a significant incident. Rationalising IAM — enforcing SSO, retiring long-lived keys, scoping permissions down, auditing group memberships — is the highest-impact action because it reduces the blast radius of every other failure mode.

Phase 3: Mature but fragmented. This is Westmoor University and, at a different scale, Brevenmark Assurance Group. The organisation has a working cloud security programme for some of its estate. The programme is mature in the parts it covers. The problem is coverage: significant portions of the cloud estate sit outside the programme because they were procured by different teams, run by different functions, or arrived through acquisition. The central programme has good controls; the federated reality has partial controls or no controls, depending on how far from the centre you look.

The first priority in Phase 3 is perimeter: a credible inventory of everything that is actually in the cloud estate, followed by a mechanism for extending the programme's controls to the portions that are currently outside it. This is harder than it sounds in federated organisations, because extending the programme often requires negotiating with teams that have never been subject to it and that have their own views about what security means in their context. The chapters on multi-cloud coordination, hybrid environments, and governance return to this repeatedly. The short version: the programme that works for the central IT team but not for the research faculties or the acquired subsidiaries is not a programme. It is a programme with a large number of known exceptions, and known exceptions are where the incidents accumulate.

1.0.4 What cloud security is not

Drawing the boundary crisply is important because cloud security is a term that attracts adjacent concerns, and the conflation causes real harm to working programmes.

Cloud security is not the same as cloud compliance. Compliance is the demonstration that specific controls required by a specific regulatory regime are in place. It is a subset of security, and an important subset, but it is not the whole of it. An organisation can be compliant with ISO 27001, SOC 2 Type II, or PCI DSS and still have a materially misconfigured cloud estate — because compliance frameworks certify a point-in-time posture against a defined control set, and the defined control set may not capture the specific risk the organisation is running. Conversely, an organisation that has built a genuinely strong cloud security programme will find compliance easier to achieve and maintain, because the controls that make a programme strong tend to overlap substantially with the controls that compliance frameworks require. The relationship is useful — compliance as a floor, security as the target — provided neither is mistaken for the other.

Cloud security is not the same as cloud operations or site reliability engineering, though the disciplines share significant tooling and vocabulary. SRE is concerned with availability, reliability, and performance. Security is concerned with confidentiality, integrity, and the consequences of adversarial action. These are different objectives, and they sometimes point in different directions. The SRE's instinct to prioritise availability — to keep the system running, to minimise downtime, to move fast in a recovery — can create tension with the security function's instinct to preserve evidence, limit the blast radius, and treat an incident as a potential crime scene before treating it as a service interruption. Neither instinct is wrong. They need to coexist with a shared escalation path and a clear authority structure for the moments they conflict.

Cloud security is not purely a technical discipline. This bears emphasis because the cloud security industry — the conference talks, the vendor marketing, the certification curricula — strongly implies that it is. The cloud security problem is substantially a problem of organisational clarity: who owns the security posture of which cloud accounts, who has the authority to enforce controls on teams that have not accepted them, how procurement decisions get made in ways that don't create unassessable vendor exposure, and what the consequence is when a team does not meet the programme's expectations. None of

those questions have technical answers. They are governance and culture questions, and a programme that treats them as out of scope for the security function will spend its time remediating the same misconfigurations over and over because it never addressed the conditions that produce them.

Cloud security is not a one-time project. This is the trap that organisations entering Phase 1 most frequently fall into: they run a cloud security audit, close the findings, and consider the work done. Cloud environments are dynamic. They change every time a developer pushes infrastructure-as-code, every time a new service is procured, every time an identity is provisioned or an employee leaves. The controls that were correctly configured on Tuesday can drift by Friday. The asset inventory that was current last quarter may be missing three new accounts by this quarter. The cloud security programme is not a project. It is a continuous operational discipline.

1.0.5 The thesis

The thesis stated at the start of this chapter is that cloud security is a specialisation of existing security disciplines, not a replacement for them. Most of what a security professional needs in a cloud environment is what they have always needed: asset inventory, access control, logging, vulnerability management, incident response, vendor due diligence, governance and assurance. The cloud context changes where these disciplines apply, which tools implement them, and what the specific failure modes look like — but the underlying logic of the disciplines remains recognisable.

The exception list is real. There are cloud-specific failure modes that require new vocabulary and new operational practice: IAM at cloud scale, where the attack surface of an identity is not a password but a policy document that may grant permissions across thousands of services; configuration drift in ephemeral infrastructure, where the concept of “patching a server” has been replaced by “replacing an image” and the security implications of that substitution are non-trivial; the cryptographic surface of cloud key management, where the security of data at rest and in transit is now a function of how key rotation, key

access policy, and HSM hierarchy are configured; multi-cloud coordination, where the same underlying security objective requires different tooling, different data schemas, and different operational procedures for each provider in the mix. These are genuinely new material. They are also learnable material, and the rest of this book is one route through learning them.

The practitioners who do well in cloud security are not the ones who treat it as a wholly new discipline requiring new tools and new expertise from scratch. They are the ones who correctly identify what transfers from their existing practice, add the cloud-specific layer to it deliberately, and maintain the discipline of doing both continuously. A network security engineer who understands how traffic flows and where to put controls has a significant head start on cloud network security. An IAM practitioner who understands least-privilege and the abuse patterns of overpermissioned identities has a significant head start on cloud IAM. A forensics analyst who understands evidence preservation has a significant head start on cloud forensics, with the genuinely new wrinkle that many cloud artefacts exist only in flight and must be captured before the session or the instance that generated them is gone.

1.0.6 The four organisations

The book uses four recurring organisations as its primary worked examples. They are described in full in Chapter 4, but a sketch of each helps orient the reader from the start.

Stagheart Studio is a fifteen-person remote-first product-design agency in Bristol, co-founded by Esme Caldwell and Ravi Patnaik. They run on Google Workspace, Figma, GitHub, and a single AWS account hosting a portfolio site, a contact-form Lambda, and three internal microservices behind a client portal. Stagheart is the example of organisations that have adopted cloud services without a corresponding security programme and need to get to minimum-viable practice without the resources to hire a security specialist.

Helix Mobile is a six-hundred-person mobile virtual network operator in Dublin, subject to GDPR, the EU Electronic Communications

Code, and NIS2. Their multi-cloud AWS and Azure platform was built for speed and is now carrying the accumulated security debt of six years of rapid growth. Aoife Brennan, recently hired as Head of Cloud Security, is leading the programme rebuild.

Westmoor University is a research-intensive university in the north of England with eight thousand staff, twenty-six thousand students, and a genuinely heterogeneous cloud estate spanning central IT, research faculties, and federated partner institutions. Dr Catriona MacLeod is the CISO. The central programme is mature. The federated reality is not.

Brevenmark Assurance Group is a fifty-thousand-person global insurance holding company in Zürich, operating a mature multi-cloud programme across AWS, Azure, and GCP, in nine major jurisdictions with a forty-five-person Cloud Security Platform Team. The programme is strong and still not finished: two recent acquisitions sit outside the landing-zone, DORA obligations are pressing, and the M&A backlog is growing faster than the integration team can clear it.

1.0.7 The shape of the rest of the book

The book runs to eighteen chapters across three parts.

Part I — Chapters 1 through 5 — establishes the foundations: this chapter on what cloud security is and is not; Chapter 2 on the shared-responsibility model in practice; Chapter 3 on cloud asset inventory; Chapter 4 on the four-tier maturity spectrum; and Chapter 5 on how cloud security differs from on-premises practice.

Part II — Chapters 6 through 12 — covers the control domains: IAM at cloud scale, secrets and the cryptographic surface, network and data-plane security, configuration baselines and drift, logging and SIEM integration, data security, and workload protection.

Part III — Chapters 13 through 18 — covers the platform layer: Kubernetes and container orchestration, infrastructure as code, SaaS security posture, cloud cryptography patterns, the edge, and a closing chapter on the secured platform and what comes next.

This book is complete on its own: the foundations, the control domains, and the platform layer are the full security architecture. Two

companion books extend the practice for readers who want them — *Operations and Incident Response* runs the architecture under pressure, and the *Cloud Security Operator's Workbook* collects the templates, checklists, decision trees, and playbooks in a form a working team can use directly. Each stands alone; none assumes you have read the others.

1.0.8 A note on provider agnosticism

This book is written to be useful regardless of which cloud provider or combination of providers the reader's organisation uses. Where specific services or features are cited, they are cited because the concept is well illustrated by a provider-specific example, not because that provider is the recommended or only option. AWS, Azure, and GCP are all represented across the worked examples. The security principles that govern cloud programmes — least privilege, defence in depth, configuration-as-code, continuous monitoring, thorough logging — apply regardless of provider. The tooling is provider-specific; the principles are not.

1.0.9 A note on defensibility

This book offers practitioner framings drawn from real cloud security operations. The framings are intended to support security professionals in thinking clearly about cloud environments and in extending their existing craft to defend them. They are not legal advice, regulatory determinations, or guarantees of any particular cloud security outcome. The cloud security field is moving quickly, and the regulatory environment — NIS2, DORA, the EU AI Act, US federal rulemaking, and sector-specific requirements across financial services, health-care, and critical infrastructure — is in active flux. The reader carries the responsibility of applying these framings to their own context and consulting qualified counsel on questions of obligation.

What the book does claim — and the claim is bounded — is that the disciplines it describes are recognisable as security disciplines, have

been practised in production cloud environments by working security functions, and that an organisation that adopts them in a form appropriate to its scale and risk appetite will be in a meaningfully better position to defend its cloud workloads than one that does not.

Recap: cloud security is the application of security disciplines to workloads, data, identities, and services operated on third-party cloud infrastructure; the shared-responsibility model defines the split between what the provider secures and what the customer secures; most cloud security incidents originate in the customer's layer; and the book's four case-study organisations — Stagheart, Helix, Westmoor, and Brevenmark — represent the range of scales and contexts in which cloud security problems appear.

Next: Chapter 2 — The Shared-Responsibility Model in Practice. Before a cloud security programme can be designed, the organisation must understand precisely what it owns. Chapter 2 walks the shared-responsibility model across IaaS, PaaS, and SaaS service types, shows what “customer responsibility” translates to in operational terms for each, and names the misconfiguration categories that arise most predictably when organisations underestimate the scope of what they own.

You've just read Chapter 1 of

Cloud Security

The full book runs 326 pages across 18 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/cloud-security

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



sylvanassurance.com/go/free-cloud-security

*This sample is provided for evaluation. The full text is © 2026 Sylvan Assurance, LLC.
This material is provided for general information and does not constitute legal advice.*